

Stability of ε -Kernels

Pankaj K. Agarwal

Jeff M. Phillips

Hai Yu

March 31, 2010

Abstract

Given a set P of n points in \mathbb{R}^d , an ε -kernel $K \subseteq P$ approximates the directional width of P in every direction within a relative $(1 - \varepsilon)$ factor. In this paper we study the stability of ε -kernels under dynamic insertion and deletion of points to P and by changing the approximation factor ε . In the first case, we say an algorithm for dynamically maintaining a ε -kernel is stable if at most $O(1)$ points change in K as one point is inserted or deleted from P . We describe an algorithm to maintain an ε -kernel of size $O(1/\varepsilon^{(d-1)/2})$ in $O(1/\varepsilon^{(d-1)/2} + \log n)$ time per update. Not only does our algorithm maintain a stable ε -kernel, its update time is faster than any known algorithm that maintains an ε -kernel of size $O(1/\varepsilon^{(d-1)/2})$. Next, we show that if there is an ε -kernel of P of size κ , which may be dramatically less than $O(1/\varepsilon^{(d-1)/2})$, then there is an $(\varepsilon/2)$ -kernel of P of size $O(\min\{1/\varepsilon^{(d-1)/2}, \kappa^{\lfloor d/2 \rfloor} \log^{d-2}(1/\varepsilon)\})$. Moreover, there exists a point set P in \mathbb{R}^d and a parameter $\varepsilon > 0$ such that if every ε -kernel of P has size at least κ , then any $(\varepsilon/2)$ -kernel of P has size $\Omega(\kappa^{\lfloor d/2 \rfloor})$.¹

¹Research supported by subaward CIF-32 from NSF grant 0937060 to CRA, by NSF under grants CNS-05-40347, CFF-06-35000, and DEB-04-25465, by ARO grants W911NF-04-1-0278 and W911NF-07-1-0376, by an NIH grant 1P50-GM-08183-01, by a DOE grant OEG-P200A070505, and by a grant from the U.S.–Israel Binational Science Foundation.

1 Introduction

With recent advances in sensing technology, massive geospatial data sets are being acquired at an unprecedented rate in many application areas, including GIS, sensor networks, robotics, and spatial databases. Realizing the full potential of these data sets requires developing scalable algorithms for analyzing and querying them. Among many interesting algorithmic developments to meet this challenge, there is an extensive amount of work on computing a “small summary” of large data sets that preserves certain desired properties of the input data and on obtaining a good trade-off between the quality of the summary and its size. A coreset is one example of such approximate summaries. Specifically, for an input set P and a function f , a *coreset* $C \subseteq P$ is a subset of P (with respect to f) with the property that $f(C)$ approximates $f(P)$. If a small-size coreset C can be computed quickly (much faster than computing $f(P)$), then one can compute an approximate value of $f(P)$ by first computing C and then computing $f(C)$. This coreset-based approach has been successfully used in a wide range of geometric optimization problems over the last decade. See [2] for a survey.

ε -kernels. Agarwal *et al.* [1] introduced the notion of ε -kernels and proved that it is a coreset for many functions. For any direction $u \in \mathbb{S}^{d-1}$, let $P[u] = \arg \max_{p \in P} \langle p, u \rangle$ be the extreme point in P along u ; $\omega(P, u) = \langle P[u] - P[-u], u \rangle$ is called the *directional width* of P in direction u . For a given $\varepsilon > 0$, $K \subset P \subset \mathbb{R}^d$ is called an ε -kernel of P if

$$\langle P[u] - K[u], u \rangle \leq \varepsilon \omega(P, u)$$

for all directions $u \in \mathbb{S}^{d-1}$.² For simplicity, we assume $\varepsilon \in (0, 1)$, because for $\varepsilon \geq 1$, one can choose a constant number of points to form an ε -kernel, and we assume d is constant. By definition, if X is an ε -kernel of P and K is a δ -kernel of X , then K is a $(\delta + \varepsilon)$ -kernel of P .

Agarwal *et al.* [1] showed that there exists an ε -kernel of size $O(1/\varepsilon^{(d-1)/2})$ and it can be computed in time $O(n + 1/\varepsilon^{3d/2})$. The running time was improved by Chan [6] to $O(n + 1/\varepsilon^{d-3/2})$ (see also [12]). In a number of applications, the input point set is being updated periodically, so algorithms have also been developed to maintain ε -kernels dynamically. Agarwal *et al.* [1] had described a data structure to maintain an ε -kernel of size $O(1/\varepsilon^{(d-1)/2})$ in $(\log(n)/\varepsilon)^{O(d)}$ time per update. The update time was recently improved by Chan [7] to $O((1/\varepsilon^{(d-1)/2}) \log n + 1/\varepsilon^{d-3/2})$. His approach can also maintain an ε -kernel of size $O((1/\varepsilon^d) \log n)$ with update time $O(\log n)$. If only insertions are allowed (e.g. in a streaming model), the size of the data structure can be improved to $O(1/\varepsilon^{(d-1)/2})$ [3, 13].

In this paper we study two problems related to the *stability* of ε -kernels: how ε -kernels change as we update the input set or vary the value of ε .

Dynamic stability. Since the aforementioned dynamic algorithms for maintaining an ε -kernel focus on minimizing the size of the kernel, changing a single point in the input set P may drastically change the resulting kernel. This is particularly undesirable when the resulting kernel is used to build a dynamic data structure for maintaining another information. For example, kinetic data structures (KDS) based on coresets have been proposed to maintain various extent measures of a set of moving points [2]. If an insertion or deletion of an object changes the entire summary, then one has to reconstruct the entire KDS instead of locally updating it. In fact, many other dynamic data structures for maintaining geometric summaries also suffer from this undesirable property [4, 9, 11].

We call an ε -kernel *s-stable* if the insertion or deletion of a point causes the ε -kernel to change by at most s points. For brevity, if $s = O(1)$, we call the ε -kernel to be *stable*. Chan’s dynamic algorithm can

²This is a slightly stronger version of the definition than defined in [1] and an ε -kernel K gives a relative $(1 + 2\varepsilon)$ -approximation of $\omega(P, u)$ for all $u \in \mathbb{S}^{d-1}$ (i.e. $\omega(K, u) \leq \omega(P, u) \leq (1 + 2\varepsilon)\omega(K, u)$).

be adapted to maintain a stable ε -kernel of size $O((1/\varepsilon^{d-1}) \log n)$; see Lemma 2.6 below. An interesting question is whether there is an efficient algorithm for maintaining a stable ε -kernel of size $O(1/\varepsilon^{(d-1)/2})$, as points are being inserted or deleted. Maintaining a stable ε -kernel dynamically is difficult for two main reasons. First, for an input set P , many algorithms compute ε -kernels in two or more steps. They first construct a large ε -kernel K' (e.g. see [1, 7]), and then use a more expensive algorithm to create a small ε -kernel of K' . However, if the first algorithm is unstable, then K' may change completely each time P is updated. Second, all of the known ε -kernel algorithms rely on first finding a “rough shape” of the input set P (e.g., finding a small box that contains P), estimating its fatness [5]. This rough approximation is used crucially in the computation of the ε -kernel. However, this shape is itself very unstable under insertions or deletions to P . Overcoming these difficulties, we prove the following in Section 2:

Theorem 1.1. *Given a parameter $0 \leq \varepsilon \leq 1$, a stable ε -kernel of size $O(1/\varepsilon^{(d-1)/2})$ of a set of n points in \mathbb{R}^d can be maintained under insertions and deletions in $O(1/\varepsilon^{(d-1)/2} + \log n)$ time.*

Note that the update time of maintaining an ε -kernel of size $O(1/\varepsilon^{(d-1)/2})$ is better than that in [7].

Approximation stability. If the size of an ε -kernel K is $O(1/\varepsilon^{(d-1)/2})$, then decreasing ε changes K quite predictably. However, this is the worst-case bound, and it is possible that the size of K may be quite small, e.g., $O(1)$, or in general much smaller than the $1/\varepsilon^{(d-1)/2}$ maximum (efficient algorithms are known for computing ε -kernels of near-optimal size [2]). Then how much can the size increase as we reduce the allowable error from ε to $\varepsilon/2$? For any $\varepsilon > 0$, let $\kappa(P, \varepsilon)$ denote the minimum size of an ε -kernel of P . Unlike many shape simplification problems, in which the size of simplification can change drastically as we reduce the value of ε , we show (Section 3) that this does not happen for ε -kernels and that $\kappa(P, \varepsilon/2)$ can be expressed in terms of $\kappa(P, \varepsilon)$.

Theorem 1.2. *For any point set P and for any $\varepsilon > 0$,*

$$\kappa(P, \varepsilon/2) = O(\min\{\kappa(P, \varepsilon)^{\lfloor d/2 \rfloor} \log^{d-2}(1/\varepsilon), 1/\varepsilon^{(d-1)/2}\}).$$

Moreover, there exist a point set P and some $\varepsilon > 0$ such that $\kappa(P, \varepsilon/2) = \Omega(\kappa(P, \varepsilon)^{\lfloor d/2 \rfloor})$.

2 Dynamic Stability

In this section we describe an algorithm that proves Theorem 1.1. The algorithm is composed of a sequence of modules, each with certain property. We first define the notion of anchor points and fatness of a point set and describe two algorithms for maintaining stable ε -kernels with respect to a fixed anchor: one of them maintains a kernel of size $O(1/\varepsilon^{d-1})$ and the other of size $O(1/\varepsilon^{(d-1)/2})$; the former has smaller update time. Next, we briefly sketch how Chan’s algorithm [7] can be adapted to maintain a stable ε -kernel of size $O(1/\varepsilon^{(d-1)} \log n)$. Then we describe the algorithm for updating anchor points and maintaining a stable kernel as the anchors change. Finally, we put these modules together to obtain the final algorithm. We make the following simple observation, which will be crucial for combining different modules.

Lemma 2.1 (Composition Lemma). *If K is an s -stable ε -kernel of P and K' is an s' -stable ε' -kernel of K , then K' is an $(s \cdot s')$ -stable $(\varepsilon + \varepsilon')$ -kernel of P .*

Anchors and fatness of a point set. We call a point set P β -fat if

$$\max_{u, v \in \mathbb{S}^{d-1}} \omega(P, u) / \omega(P, v) \leq \beta.$$

If β is a constant, we sometimes just say that P is *fat*. An arbitrary point set P can be made fat by applying an affine transform: we first choose a set of $d + 1$ *anchor points* $A = \{a_0, a_1, \dots, a_d\}$ using the following procedure of Barequet and Har-Peled [5]. Choose a_0 arbitrarily. Let a_1 be the farthest point from a_0 . Then inductively, let a_i be the farthest point from the flat $\text{span}(a_0, \dots, a_{i-1})$. (See Figure 1.) The anchor points A define a bounding box I_A with center at a_0 and orthogonal directions defined by vectors from the flat $\text{span}(a_0, \dots, a_{i-1})$ to a_i . The extents of I_A in each orthogonal direction is defined by placing each a_i on a bounding face and extending I_A the same distance from a_0 in the opposite direction. Next we perform an affine transform T_A on P such that the vector from the flat $\text{span}(a_0, \dots, a_{i-1})$ to a_i is equal to e_i , where $e_0 = (0, \dots, 0)$, $e_1 = (1, 0, \dots, 0)$, \dots , $e_d = (0, \dots, 0, 1)$. This ensures that $T_A(P) \subseteq T_A(I_A) = [-1, 1]^d$. The next lemma shows that $T_A(P)$ is fat.

Lemma 2.2. *For all $u \in \mathbb{S}^{d-1}$ and for $\beta_d \leq 2^d d^{3/2} d!$,*

$$\omega(T_A(A), u) \leq \omega(T_A(P), u) \leq \omega(T_A(I_A), u) \leq \beta_d \cdot \omega(T_A(A), u). \quad (2.1)$$

Proof. The first two inequalities follow by $A \subset P \subset I_A$. We can upper bound $\max_u \omega(T_A(I_A), u) = \omega([-1, 1]^d, u) \leq \sqrt{d}$. The volume of the convex hull $\text{conv}(T_A(A))$ is $1/d!$ since it is a d -simplex and $\langle T_A(a_0) - T_A(a_i), e_i \rangle = 1$ for each direction e_i . We can then scale $\text{conv}(T_A(A))$ by a factor $1/2$ (shrinking the volume by factor $1/2^d$) so it fits in $[0, 1]^d$. Now we can apply a lemma from [8] that the minimum width of a convex shape that is contained in $[0, 1]^d$ is at least $1/d$ times its d -dimensional volume, which is $1/(2^d d!)$. The fatness of $T_A(P)$ follows from the fatness of $T_A(I_A)$. \square




Figure 1. Anchor points $A = \{a_0, a_1, a_2\}$, rectangle I_A , and transform T_A applied to P ; square \mathbb{H} , two-dimensional grid \mathbb{G} , and one-dimensional grid \mathbb{G}_e on the edge e of \mathbb{H} .

Agarwal *et al.* [1] show if K is an ε -kernel of P , then $T(K)$ is an ε -kernel of $T(P)$ for any affine transform T , which implies that one can compute an ε -kernel of $T(P)$. We will need the following generalization of the definition of ε -kernel. For two points sets P and Q , a subset $K \subseteq P$ is called an ε -kernel of P with respect to Q if $\langle P[u] - K[u], u \rangle \leq \varepsilon \omega(Q, u)$ for all $u \in \mathbb{S}^{d-1}$.

Stable ε -kernels for a fixed anchor. Let A be a set of anchor points of P , as described above. We describe algorithms for maintaining stable ε -kernels (with respect to A) under the assumption that A remains a set of anchor points of P , i.e., $A \subseteq P \subset I_A$, as P is being updated by inserting and deleting points. In view of the above discussion, without loss of generality, we assume $I_A = [-1, +1]^d$ and denote it by \mathbb{H} . As for the static case [1, 6], we first describe a simpler algorithm that maintains a stable ε -kernel of size $O(1/\varepsilon^{d-1})$, and then a more involved one that maintains a stable ε -kernel of size $O(1/\varepsilon^{(d-1)/2})$.

Set $\delta = \varepsilon/\sqrt{d}$ and draw a d -dimensional grid \mathbb{G} inside \mathbb{H} of size δ , i.e., the side-length of each grid cell is at most δ ; \mathbb{G} has $O(1/\delta^d)$ cells. For each grid cell τ , let $P_\tau = P \cap \tau$. For a point $x \in \mathbb{H}$ lying in a grid cell τ , let \hat{x} be the vertex of τ nearest to the origin; we can view x being *snapped* to the vertex \hat{x} . For each facet f of \mathbb{H} , \mathbb{G} induces a $(d-1)$ -dimensional grid \mathbb{G}_f on f ; \mathbb{G} contains a *column* of cells for each cell in \mathbb{G}_f . For each cell $\Delta \in \mathbb{G}_f$, we choose (at most) one point of P as follows: let τ be the nonempty grid cell in the column of \mathbb{G} corresponding to Δ that is closest to f . We choose an arbitrary point from P_τ ; if there is

no nonempty cell in the column, no point is chosen. Let L_f be the set of chosen points. Set $\mathcal{L} = \bigcup_{f \in \mathbb{H}} L_f$. Agarwal *et al.* [1] proved that \mathcal{L} is an ε -kernel of P . Insertion or deletion of a point in P affects at most one point in L_f , and it can be updated in $O(\log(1/\varepsilon))$ time. Hence, we obtain the following:

Lemma 2.3. *Let P be a set of n points in \mathbb{R}^d , let $A \subseteq P$ be a set of anchor points of P , and let $0 < \varepsilon < 1$ be a parameter. P can be preprocessed in $O(n + 1/\varepsilon^{d-1})$ time, so that a $(2d)$ -stable ε -kernel of P with respect to A of size $O(1/\varepsilon^{d-1})$ can be maintained in $O(\log 1/\varepsilon)$ time per update provided that A remains an anchor set of P .*

Agarwal *et al.* [1] and Chan [6] have described algorithms for computing an ε -kernel of size $O(1/\varepsilon^{(d-1)/2})$. We adapt Chan's algorithm to maintain a stable ε -kernel with respect to a fixed anchor A . We begin by mentioning a result of Chan that lies at the heart of his algorithm.

Lemma 2.4 (Chan [6]). *Let $E \in \mathbb{N}$, $E^\tau \leq F \leq E$ for some $0 < \tau < 1$, and $P \subseteq [0 : E]^{d-1} \times \mathbb{R}$ a set of at most n points. For all grid points $b \in [0 : F]^{d-1} \times \mathbb{R}$, the nearest neighbors of each b in P can be computed in time $O(n + E^{d-2}F)$.*

We now set $\gamma = \sqrt{\varepsilon}/c$ for a constant $c > 1$ to be used in a much sparser grid than with δ . Let $\mathbb{C} = [-2, +2]^d$ and f be a facet of \mathbb{C} . We draw a $(d-1)$ -dimensional grid on f of size γ . Assuming f lies on the plane $x_d = -2$, we choose a set $B_f = \{(i_1\gamma, \dots, i_{d-1}\gamma, -2) \in \mathbb{Z}^d \mid -\lceil 2/\gamma \rceil \leq i_1, \dots, i_{d-1} \leq \lceil 2/\gamma \rceil\}$ of grid points. For a subset $X \subseteq P$ and a point b , we define $\psi(X, b) = \arg \min_{x \in X} \|\hat{x} - b\|$, i.e., the point in X such that the snapped point is nearest to b . For a set R , $\psi(X, R) = \{\psi(X, r) \mid r \in R\}$. There is a one to one mapping between the faces of \mathbb{C} and \mathbb{H} , so we also use f to denote the corresponding facet of \mathbb{H} . Let L_f be the set of points chosen in the previous algorithm corresponding to facet f of \mathbb{H} for computing an $(\varepsilon/2)$ -kernel of P . Set $G_f = \psi(L_f, B_f)$. Chan showed that $\mathcal{G} = \bigcup_{f \in \mathbb{C}} G_f$ is an $(\varepsilon/2)$ -kernel of \mathcal{L} and thus an ε -kernel of P . Scaling \mathbb{G} and B_f appropriately and using Lemma 2.4, G_f can be computed in $O(n + 1/\varepsilon^{d-3/2})$ time. Hence, \mathcal{G} can be computed in $O(n + 1/\varepsilon^{d-3/2})$ time.

Note that $\psi(L_f, b)$ can be the same for many points $b \in B_f$, so insertion or deletion of a point in P (and thus in L_f) may change G_f significantly, thereby making \mathcal{G} unstable. We circumvent this problem by introducing two new ideas. First, $\psi(L_f, B_f)$ is computed in two stages, and second it is computed in an iterative manner. We describe the construction and the update algorithm for f ; the same algorithm is repeated for all facets.

We partition \mathbb{H} into $O(1/\gamma^{d-1})$ boxes: for $J = \langle i_1, \dots, i_{d-1} \rangle \in [-1/\gamma, 1/\gamma]^{d-1} \cap \mathbb{Z}^{d-1}$, we define $\mathbb{H}_J = [i_1\gamma, (i_1+1)\gamma] \times \dots \times [i_{d-1}\gamma, (i_{d-1}+1)\gamma] \times [-1, +1]$. We maintain a subset $X \subseteq L_f$. Initially, we set $X = L_f$. Set $X_J = X \cap \mathbb{H}_J$. We define a total order on the points of B_f . Initially, we sort B_f in lexicographic order, but the ordering will change as insertions and deletions are performed on P . Let $\langle b_1, \dots, b_u \rangle$ be the current ordering of B_f . We define a map $\varphi : B_f \rightarrow L_f$ as follows. Suppose $\varphi(b_1), \dots, \varphi(b_{i-1})$ have been defined. Let $J_i = \arg \min_J \|\hat{\psi}(X_J, b_i) - b_i\|$; here $\hat{\psi}(\cdot)$ denotes the snapped point of $\psi(\cdot)$. We set $\varphi(b_i) = \psi(X_{J_i}, b_i)$. We delete $\varphi(b_i)$ from X (and from X_{J_i}) and recompute $\hat{\psi}(X_{J_i}, B_f)$. Set $K_f = \{\varphi(b) \mid b \in B_f\}$ and $\mathcal{K} = \bigcup_f K_f$. Computing J_i and $\varphi(b_i)$ takes $O(1/\varepsilon^{(d-1)/2})$ time, and, by Lemma 2.4, $\psi(X_{J_i}, B_f)$ can be computed in $O(|X_J| + 1/\gamma^{d-2} \cdot 1/\gamma) = O(1/\varepsilon^{(d-1)/2})$ time.

It can be proved that the map φ and the set K_f satisfy the following properties:

- (P1) $\varphi(b_i) \neq \varphi(b_j)$ for $i \neq j$,
- (P2) $\varphi(b_i) = \psi(L_f \setminus \{\varphi(b_j) \mid j < i\}, b_i)$,
- (P3) $K_f \supseteq \psi(L_f, B_f)$.

Indeed, (P1) and (P2) follow from the construction, and (P3) follows from (P2). (P3) immediately implies that \mathcal{K} is an ε -kernel of P . Next, we describe the procedures for updating K_f when L_f changes. These procedures maintain (P1)–(P3), thereby ensuring that the algorithm maintains an ε -kernel.

Inserting a point. Suppose a point p is inserted into L_f . We add p to X . Suppose $p \in \mathbb{H}_J$. We recompute $\psi(X_J, B_f)$. Next, we update $\varphi(\cdot)$ and \mathcal{K} as follows. We maintain a point $\xi \in L_f$. Initially, ξ is set to p . Suppose we have processed b_1, \dots, b_{i-1} . Let $\eta \in L_f$ be the current $\varphi(b_i)$. If $\|\hat{\xi} - b_i\| \leq \|\hat{\eta} - b_i\|$, then we swap ξ and $\varphi(b_i)$, otherwise neither ξ nor $\varphi(b_i)$ is updated. We then process b_{i+1} . After processing all points of B_f if $\xi = p$, i.e., no $\varphi(b_i)$ is updated, we stop. Otherwise, we add p to K_f and delete ξ from K_f . The insertion procedure makes at most two changes in K_f , and it can be verified that (P1)-(P3) are maintained.

Deleting a point. Suppose p is deleted from L_f . Suppose $p \in \mathbb{H}_J$. If $p \notin K_f$, then $p \in X$. We delete p from X and X_J and recompute $\psi(X_J, B)$. If $p \in K_f$, i.e., there is a $b_i \in B$ with $p = \varphi(b_i)$, then $p \notin X$. We delete p from K_f and \mathcal{K} , recompute $\varphi(b_i)$, and add the new $\varphi(b_i)$ to K_f . Let $\varphi(b_i) \in \mathbb{H}_J$; we remove $\varphi(b_i)$ from X_J and recompute $\psi(X_J, B_f)$. We modify the ordering of B_f by moving b_i from its current position to the end. This is the only place where the ordering of B_f is modified. Since b_i is now the last point in the ordering of B_f , the new $\varphi(b_i)$ does not affect any other $\varphi(b_j)$. The deletion procedure also makes at most two changes in K_f and maintains (P1)-(P3).

Finally, insertion or deletion of a point in P causes at most one insertion plus one deletion in L_f , therefore we can conclude the following:

Lemma 2.5. *Let P be a set of n points in \mathbb{R}^d , A a set of anchor points of P , and $0 < \varepsilon < 1$ a parameter. P can be preprocessed in $O(n + 1/\varepsilon^{d-1})$ time into a data structure so that a stable ε -kernel of P with respect to A of size $O(1/\varepsilon^{(d-1)/2})$ can be maintained in $O(1/\varepsilon^{(d-1)/2})$ time under insertion and deletion, provided that A remains an anchor set of P .*

Stabilizing Chan's dynamic algorithm. We now briefly describe how Chan's [7] dynamic ε -kernel algorithm can be adapted so that it maintains a stable ε -kernel of size $O((1/\varepsilon^{d-1}) \log n)$. He bypasses the need of fixed anchors by partitioning P into $h = O(\log n)$ layers $\langle P_1, \dots, P_h \rangle$, where P_1 is the inner-most layer and P_h is the outer-most layer, $|P_i| \geq \gamma \sum_{j=i+1}^h |P_j|$ for a constant $\gamma > 1$, and $|P_h| = 1/\varepsilon^{d-1}$. P_1 is constructed first, and then the rest of the layers are constructed recursively with the remaining points. For each set P_i (for $i < h$) there exists a set of points A_i which serve as anchor points for P_i in the sense that they define a bounding box I_{A_i} (i.e., $P_i \subset I_{A_i}$). Furthermore, for all $u \in \mathbb{S}^{d-1}$ we have $\omega(A_i, u) \leq \omega(P_{i+1}, u)$, and this remains true for $\alpha|P_i|$ insertions or deletions to P for a constant $0 \leq \alpha \leq 1$. After $\alpha|P_i|$ updates in P , the layers P_i, P_{i+1}, \dots, P_h are reconstructed. Also at this point layers $P_j, P_{j+1}, \dots, P_{i-1}$ will need to be reconstructed if layer P_j is scheduled to be reconstructed in fewer than $\alpha|P_j|$ updates. We set $K_h = P_h$, and for $i < h$, an ε -kernel K_i of P_i with respect to A_i is maintained using Lemma 2.3. The set $\mathcal{K} = \bigcup_{i=1}^h K_i$ is an ε -kernel of P ; $|\mathcal{K}| = O((1/\varepsilon^{d-1}) \log n)$.

When a new point p is inserted into P , it is added to the outermost layer P_i , i.e., i is the largest such value, such that $p \in I_{A_i}$. If a point is inserted into or deleted from P_i , we update K_i using Lemma 2.3. The update time follows from the following lemma.

Lemma 2.6. *For any $0 < \varepsilon < 1$, an ε -kernel \mathcal{K} of P of size $O((1/\varepsilon^{d-1}) \log n)$ can be maintained in $O(\log n)$ time, and the number of changes in \mathcal{K} at each update is $O(1)$.*

Proof. Recall that if the insertion or deletion of a point does not require reconstruction of any layers, the update time is $O(\log n)$ and by Lemma 2.3, only $O(1)$ changes occur in \mathcal{K} . We start by bounding the amortized time spent in reconstructing layers and their kernels.

The kernel K_i of P_i is rebuilt if at least $\alpha|P_i|$ updates have taken place since the last reconstruction, if K_l needs to be rebuilt for some $l < i$, or if some K_j is rebuilt for $j > i$ and K_i is scheduled to be rebuilt in fewer than $\alpha|P_j|$ updates to P . These conditions imply that the i th layer is rebuilt after every $\alpha|P_h|2^j$ updates where j is the smallest integer such that $|P_i| \leq |P_h|2^j$. The third condition acts to coordinate the

rebuilding of the layers so that the i th layer is not rebuilt after fewer than $\alpha|P_i|/2$ updates since its last rebuild.

K_h is rebuilt after α/ε^{d-1} updates. And $|P_i| \leq \gamma^{h-i}|P_h|$. Since the entire system is rebuilt after $\Theta(|P_1|) = \Theta(n)$ updates, we call this interval a round. We can bound the updates to \mathcal{K} in a round by charging $O(1/\varepsilon^{d-1})$ each time a K_i is rebuilt, which occurs at most $\Theta(|P_1|)2/\alpha|P_i| = \Theta(n/|P_i|)$ times in a round.

$$\sum_{i=1}^h \frac{O(n)}{\Omega(|P_i|)} \cdot O(1/\varepsilon^{d-1}) \leq \sum_{i=1}^h \frac{O(n) \cdot O(1/\varepsilon^{d-1})}{\Omega(\gamma^{h-i}|P_h|)} = \sum_{i=1}^h \frac{O(n)}{\Omega(\gamma^i)} = O(n).$$

Thus there are $O(n)$ updates to the ε -kernel \mathcal{K} for every $\Theta(n)$ updates to P . Thus, in an amortized sense, for each update to P there are $O(1)$ updates to \mathcal{K} .

This process can be de-amortized by adapting the standard techniques for de-amortizing the update time of a dynamic data structure [10]. If a kernel K_i is valid for k insertions or deletions to P , then we start construction on the next kernel K_i after $k/2$ insertions or deletions have taken place since the last time K_i was rebuilt. All insertions can be put in a queue and added to K by the time $k/2$ steps have transpired. All deletions from old K_i to new K_i are then queued and removed from K before another $k/2$ insertions or deletions. This can be done by performing $O(1)$ queued insertions or deletions from K each insertion or deletion from P . \square

Updating anchors. We now describe the algorithm for maintaining a stable ε -kernel when anchors of P are no longer fixed and need to be updated dynamically. Roughly speaking, we divide P into *inner* and *outer* subsets of points. The outer subset acts as a *shield* so that a stable kernel of the inner subset with respect to a fixed anchor can be maintained using Lemma 2.3 or 2.5. When the outer subset can no longer act as a shield, we reconstruct the inner and outer sets and start the algorithm again. We refer to the duration between two consecutive reconstruction steps as an *epoch*. The algorithm maintains a stable kernel within each epoch, and the amortized number of changes in the kernel because of reconstruction at the beginning of a new epoch will be $O(1)$. As above, we use the same de-amortization technique to make the ε -kernel stable across epochs. We now describe the algorithm in detail.

In the beginning of each epoch, we perform the following preprocessing. Set $\alpha = 1/10$ and compute a α -kernel \mathcal{L} of P of size $O(\log n)$ using Chan's dynamic algorithm; we do not need the stable version of his algorithm described above. \mathcal{L} can be updated in $O(\log n)$ time per insertion/deletion. We choose a parameter m , which is set to $1/\varepsilon^{d-1}$ or $1/\varepsilon^{(d-1)/2}$. We create the outer subset of P by peeling off m "layers" of anchor points A_1, \dots, A_m . Initially, we set $P_0 = P$. Suppose we have constructed A_0, \dots, A_{i-1} . Set $P_{i-1} = P \setminus \bigcup_{j=1}^{i-1} A_j$, and \mathcal{L} is an α -kernel of P_{i-1} . Next, we construct the anchor set A_i of \mathcal{L} as described earlier in this section. We set $P_i = P_{i-1} \setminus A_i$ and update \mathcal{L} so that it is an α -kernel of P_i . Let $\mathcal{A} = \bigcup_i A_i$, $A = A_m$, and $P_I = P \setminus \mathcal{A}$. Let $\mathbb{H} = (1 + \alpha)I_A$. By construction $P_I \subset \mathbb{H}$. \mathcal{A} forms the outer subset and acts as a shield for P_I , which is the inner subset. Set $\delta = \varepsilon/(2(1 + \alpha)(\beta_d)^2)$, where β_d is the constant in Lemma 2.2.

If $m = 1/\varepsilon^{d-1}$ (resp. $1/\varepsilon^{(d-1)/2}$), we maintain a stable δ -kernel \mathcal{K}_I of P_I with respect to A of size $O(m)$ using Lemma 2.3 (resp. Lemma 2.5). Set $\mathcal{K} = \mathcal{K}_I \cup \mathcal{A}$; $|\mathcal{K}| = O(m)$. We prove below that \mathcal{K} is an ε -kernel of P . Let p be a point that is inserted into or deleted from P . If $p \in \mathbb{H}$, then we update \mathcal{K}_I using Lemma 2.3 or 2.5. On the other hand, if p lies outside \mathbb{H} , we insert it into or delete it from \mathcal{A} . Once \mathcal{A} has been updated m times, we end the current epoch and discard the current \mathcal{K} . We begin a new epoch and reconstruct \mathcal{A} , P_I , and \mathcal{K}_I as described above.

The preprocessing step at the beginning of a new epoch causes $O(m)$ changes in \mathcal{K} and there are at least m updates in each epoch, therefore the algorithm maintains a stable kernel in the amortized sense. As above, using a de-amortization technique, we can ensure that \mathcal{K} is stable. The correctness of the algorithm follows from the following lemma.

Lemma 2.7. \mathcal{K} is always an ε -kernel of P .

Proof. It suffices to prove the lemma for a single epoch. Since we begin a new epoch after m updates in \mathcal{A} , there is at least one i such that $A_i \subseteq \mathcal{A} \subseteq \mathcal{K}$. Thus we can show that $A_i \cup \mathcal{K}_I$ forms an ε -kernel of $A_i \cup P_I$. For any direction $u \in \mathbb{S}^{d-1}$

$$\begin{aligned}
\omega(\mathcal{K}_I, u) \leq \omega(P_I, u) &\leq \omega(\mathcal{K}_I, u) + 2(1 + \alpha)\delta \cdot \omega(A, u) \\
&\leq \omega(\mathcal{K}_I, u) + 2(1 + \alpha)\delta\beta_d \cdot \omega(P_I, u) \quad [\text{via Lemma 2.2}] \\
&\leq \omega(\mathcal{K}_I, u) + 2(1 + \alpha)\delta\beta_d \cdot \omega(P_i, u) \\
&\leq \omega(\mathcal{K}_I, u) + 2(1 + \alpha)\delta\beta_d \cdot \omega(I_i, u) \\
&\leq \omega(\mathcal{K}_I, u) + 2(1 + \alpha)\delta(\beta_d)^2 \cdot \omega(A_i, u) \quad [\text{via Lemma 2.2}] \\
&\leq \omega(\mathcal{K}_I, u) + 2(1 + \alpha)\delta(\beta_d)^2 \cdot \omega(P_I \cup A_i, u).
\end{aligned}$$

Thus, for any direction $u \in \mathbb{S}^{d-1}$ the extreme point of $P_I \cup A_i$ is either in P_I or A_i . In the first case, \mathcal{K}_I approximates the width within a factor of $2(1 + \alpha)\delta(\beta_d)^2 \cdot \omega(P_I \cup A_i, u) = \varepsilon \cdot \omega(P_I \cup A_i, u)$. In the second case, the extreme point is in \mathcal{K} because all of A_i is in \mathcal{K} . Thus the set $P_I \cup A_i$ has an ε -kernel in \mathcal{K} , and the rest of the points are also in \mathcal{K} , so \mathcal{K} is an ε -kernel of the full set P .

The size of \mathcal{K} starts at $O(m)$ because both \mathcal{K}_I and \mathcal{A} are of size $O(m)$. At most m points are inserted outside of I_A and hence into \mathcal{A} , thus the size of $\mathcal{K} = \mathcal{K}_I \cup \mathcal{A}$ is still $O(m)$ after m steps. Then the epoch ends. \square

Using Lemmas 2.3 and 2.5, we can bound the update time and conclude the following.

Lemma 2.8. For a set P of n points in \mathbb{R}^d and a parameter $0 < \varepsilon < 1$, there is a data structure that can maintain a stable ε -kernel of P of size:

- (a) $O(1/\varepsilon^{(d-1)/2})$ under insertions and deletions in time $O(n\varepsilon^{(d-1)/2} + 1/\varepsilon^{(d-1)/2} + \log n)$, or
- (b) $O(1/\varepsilon^{d-1})$ in time $O(n\varepsilon^{d-1} + \log n + \log(1/\varepsilon))$.

Proof. We build an outer kernel of size $O(m)$ in $O(n + m \log n)$ time. It lasts for $\Omega(m)$ insertions or deletions, so its construction time can be amortized over that many steps, and thus it costs $O(n/m + \log n)$ time per insertion or deletion.

In maintaining the inner kernel the preprocessing time can be amortized over m steps, but the update time cannot. In case (a) we maintain the inner kernel of size $m = O(1/\varepsilon^{(d-1)/2})$ with Lemma 2.5. The update time is $O(n\varepsilon^{(d-1)/2} + 1/\varepsilon^{(d-1)/2})$. In case (b) we maintain the inner kernel of size $m = O(1/\varepsilon^{d-1})$ with Lemma 2.3. The update time is $O(n\varepsilon^{d-1} + \log(1/\varepsilon))$.

The update time can be made worst case using a standard de-amortization techniques [10]. More specifically, we start rebuilding the inner and outer kernels after $m/2$ steps and spread out the cost over the next $m/2$ steps. We put all of the needed insertions in a queue, inserting a constant number of points to \mathcal{K} each update to P . Then after the new kernel is built, we enqueue required deletions from \mathcal{K} and perform a constant number each update to P over the next $m/2$ steps. \square

Putting it together. For a point set $P \subset \mathbb{R}^d$ of size n , we can produce the best size and update time tradeoff for stable ε -kernels by invoking Lemma 2.1 to compose three stable ε -kernel algorithms, as illustrated in Figure 2. We first apply Lemma 2.6 to maintain a stable $(\varepsilon/3)$ -kernel \mathcal{K}_1 of P of size $O(\min\{n, (1/\varepsilon^{d-1}) \log n\})$ with update time $O(\log n)$. We then apply Lemma 2.8 to maintain a stable $(\varepsilon/3)$ -kernel \mathcal{K}_2 of \mathcal{K}_1 of size $O(1/\varepsilon^{d-1})$ with update time $O(|\mathcal{K}_1|\varepsilon^{d-1} + \log |\mathcal{K}_1| + \log(1/\varepsilon)) = O(\log n + \log(1/\varepsilon))$. Finally we apply Lemma 2.8 again to maintain a stable $(\varepsilon/3)$ -kernel \mathcal{K} of \mathcal{K}_2 of size

$O(1/\varepsilon^{(d-1)/2})$ with update time $O(|\mathcal{K}_2|\varepsilon^{(d-1)/2} + 1/\varepsilon^{(d-1)/2} + \log |\mathcal{K}_2|) = O(1/\varepsilon^{(d-1)/2})$. \mathcal{K} is a stable ε -kernel of P of size $O(1/\varepsilon^{(d-1)/2})$ with update time $O(\log n + 1/\varepsilon^{(d-1)/2})$. This completes the proof of Theorem 1.1.




Figure 2. Composing stable ε -kernel algorithms.

3 Approximation Stability

In this section we prove Theorem 1.2. We first give a short proof for the lower-bound and then a more involved proof of the upper bound. For the upper bound, we first develop basic ideas and prove the theorem in \mathbb{R}^2 and \mathbb{R}^3 before generalizing to \mathbb{R}^d .

3.1 Lower Bound

Take a cyclic polytope with n vertices and $\Omega(n^{\lfloor d/2 \rfloor})$ facets and convert it into a fat polytope \mathcal{P} using standard procedures [1]. For a parameter $\varepsilon > 0$, we add, for each facet f of \mathcal{P} , a point p_f that is ε far away from the facet. Let P be the set of vertices of \mathcal{P} together with the collection of added points. We choose ε sufficiently small so that points in P are in convex position and all non-facet faces of \mathcal{P} remain as faces of $\text{conv}(P)$. Then the size of an optimal ε -kernel of P is at most n (by taking the vertices of \mathcal{P} as an ε -kernel), but the size of an optimal $\varepsilon/2$ -kernel is at least the number of facets of \mathcal{P} , because every point of the form p_f has to be present in the kernel. The first half of the lower bound is realized with $O(1/\varepsilon^{(d-1)/2})$ evenly-spaced points on a sphere, and hence the full lower bound is proved.

3.2 Upper Bound

By [1], it suffices to consider the case in which P is fat and the diameter of P is normalized to 1. Let \mathcal{K} be an ε -kernel of P of the smallest size. Let $\mathcal{P} = \text{conv}(\mathcal{K})$, and $\mathcal{P}_\varepsilon = \mathcal{P} \oplus \varepsilon \mathbb{B}^d$. We have $\mathcal{P} \subseteq \text{conv}(P) \subseteq \mathcal{P}_\varepsilon$ by the definition of ε -kernels. It suffices to show that there is a set $\mathcal{K}' \subseteq P$ such that for $\mathcal{P}' = \text{conv}(\mathcal{K}')$, $\mathcal{P}' \subseteq \text{conv}(P) \subseteq \mathcal{P}'_{\varepsilon/2}$, and $|\mathcal{K}'| = O(|\mathcal{K}|^{\lfloor d/2 \rfloor} \log^{d-2}(1/\varepsilon))$ [1].

For convenience, we assume that \mathcal{K}' is not necessarily a subset of points in P ; instead, we only require \mathcal{K}' to be a subset of points in $\text{conv}(P)$. By Caratheodory's theorem, for each point $x \in \mathcal{K}$, we can choose a set $P_x \subseteq P$ of at most $d + 1$ points such that $x \in \text{conv}(P_x)$. We set $\bigcup_{x \in \mathcal{K}'} P_x$ as the desired $(\varepsilon/2)$ -kernel of P ; $|\bigcup_{x \in \mathcal{K}'} P_x| \leq (d + 1)|\mathcal{K}'| = O(\kappa(P, \varepsilon)^{\lfloor d/2 \rfloor} \log^{d-2}(1/\varepsilon))$.

Initially, we add a point into \mathcal{K}' for each point in \mathcal{K} . If $p \in \mathcal{K}$ lies on $\partial \text{conv}(P)$, we add p to \mathcal{K}' . Otherwise we project p onto $\partial \text{conv}(P)$ in a direction in which p is maximal in \mathcal{K} and add the projected point to \mathcal{K}' . Abusing the notation slightly, we use \mathcal{P} to denote the convex hull of these initial points. For simplicity, we assume \mathcal{P} to be a simplicial polytope.

Decomposition of $\mathcal{P}_\varepsilon \setminus \text{intr } \mathcal{P}$. There are d types of simplices on $\partial \mathcal{P}$. In \mathbb{R}^2 these are points and edges. In \mathbb{R}^3 these are points, edges, and triangles. We can decompose $\mathcal{P}_\varepsilon \setminus \text{intr } \mathcal{P}$ into a set of regions, each region $\sigma(f)$ corresponding to a simplex f in \mathcal{P} . For each simplex f in \mathcal{P} let $f^* \subseteq \mathbb{S}^{d-1}$ denote the dual of f in the Gaussian diagram of \mathcal{P} . Recall that if f has dimension k ($0 \leq k \leq d - 1$), then f^* has dimension $d - 1 - k$.

The region $\mathcal{P}_\varepsilon \setminus \text{intr } \mathcal{P}$ is partitioned into a collection of $|\mathcal{P}|$ regions (where $|\mathcal{P}|$ is the number of faces of all dimensions in \mathcal{P}). Each simplex f in \mathcal{P} corresponds to a region defined

$$\sigma(f) = \{f + zu \mid 0 \leq z \leq \varepsilon, u \in f^*\}.$$

For a subsimplex $\tau \in f$, we can similarly define a region $\sigma(\tau) = \{\tau + zu \mid 0 \leq z \leq \varepsilon, u \in f^*\}$. In \mathbb{R}^2 , there are two types of regions: point regions and edge regions. In \mathbb{R}^3 , there are three types of regions: point regions (see Figure 4(a)), edge regions (see Figure 4(b)), and triangle regions (see Figure 4(c)).

For convenience, for any point $q = \bar{q} + z \cdot u \in \sigma(f)$, where $\bar{q} \in f, 0 \leq z \leq \varepsilon$, and $u \in f^*$, we write $q = \bar{q}[u, z]$ (which intuitively reads, the point whose projection onto f is \bar{q} and which is at a distance z above f in direction u). We also write $q[v] = \bar{q} + z \cdot v$ (intuitively, $q[v]$ is obtained by rotating q w.r.t. f from direction u to direction v). Similarly, we write a simplex $\bar{\Delta}[u, z] = \bar{\Delta} \oplus z \cdot u$, where $\bar{\Delta}$ is a simplex inside $f, 0 \leq z \leq \varepsilon$, and $u \in f^*$, and write $\Delta[v] = \bar{\Delta} \oplus z \cdot v$.

We will proceed to prove the upper bound as follows. For each type of region $\sigma(f)$ we place a bounded number of points from $\sigma(f) \cap \text{conv}(P)$ into \mathcal{K}' and then prove that all points in $\sigma(f) \cap \text{conv}(P)$ are within a distance $\varepsilon/2$ from some point in $\mathcal{P}' = \text{conv}(\mathcal{K}')$. We begin by introducing three ways of “gridding” $\sigma(f)$ and then use these techniques to directly prove results for several base cases, which illustrate the main conceptual ideas. These base cases will already be enough to prove the results in \mathbb{R}^2 and \mathbb{R}^3 . Finally we generalize this to \mathbb{R}^d using an involved recursive construction. We set a few global values: $\delta = \varepsilon/12d$, $\theta = 2 \arcsin(\delta/2\varepsilon)$, and $\rho = \delta/\varepsilon$.

1: Creating layers. For a point $q = \bar{q}[u, z] \in \sigma(f)$ we classify it depending on the value $z = |q - \bar{q}|$. If $z \leq \varepsilon/2$, then q is already within $\varepsilon/2$ of \mathcal{P} . We then divide the range $[\varepsilon/2, \varepsilon]$ into a constant $H = (\varepsilon/2)/\delta$ number of cases using $\mathcal{H} = \{h_1 = \varepsilon/2, h_2 = \varepsilon/2 + \delta, \dots, h_H = \varepsilon - \delta\}$. If $z \in [h_i, h_{i+1})$, then we set $q_{h_i} = \bar{q}[u, h_i]$. We define $\Psi_{f, h_i} \subset \sigma(f) \cap \text{conv}(P)$ to be the set of points that are a distance exactly h_i from f .

2: Discretize angles. We create a constant size θ -net $U_{f, h} = \{u_1, u_2, \dots\} \subset f^*$ of directions with the following properties. (1) For each $q = \bar{q}[u, h] \in \Psi_{f, h}$ there is a direction $u_i \in U_{f, h}$ such that the angle between u and u_i is at most θ . (2) For each $u_i \in U_{f, h}$ there is a point $p_i = \bar{p}_i[u_i, h] \in \Psi_{f, h}$; let $N_{f, h} = \{p_i \mid i \geq 1\}$. $U_{f, h}$ is constructed by first taking a $(\theta/2)$ -net U_f of f^* , then for each $u'_i \in U_f$ choosing a point $p_i = \bar{q}_i[u_i, h] \in \Psi_{f, h}$ where u_i is within an angle $\theta/2$ of u'_i (if one exists), and finally placing u_i in $U_{f, h}$.

3: Exponential grid. Define a set $\mathcal{D} = \{d_0, d_1 = (1 + \rho)d_0, \dots, d_m = (1 + \rho)^m d_0\}$ of distances where $d_m < 1$ and $d_0 = \delta$, so $m = O(\log 1/\varepsilon)$. For a face $f \in \mathcal{P}$, let any $r \in \sigma(f)$ be called a *support point* of f . Let p_1, \dots, p_k be the vertices of the k -simplex f . For each p_j , and each $d_i \in \mathcal{D}$ (where $d_i < \|p_j - \bar{r}\|$), let $p_{j, i}$ be the point at distance d_i from p_j on the segment $p_j \bar{r}$. For each boundary facet F of f , define a sequence of at most m simplices $F_0, F_1, \dots \in \text{conv}(F \cup \bar{r})$, each a homothet of F , so the vertices of F_i lie on segments $p_j \bar{r}$ where $p_j \in \partial F$ (see Figure 5(a)). The translation of each F_i is defined so it intersects a point $p_{j, i}$ (where $p_j \in \partial F$) and is as close to F as possible. This set of $(k - 1)$ -simplices for each F defines the exponential grid $G_{r, f}$. The full grid structure is revealed as this is applied recursively on each F_i .


The exponential grid $G_{r, \Delta}$ on a simplex Δ has two important properties for a point $q \in \Delta$:

- (G1) If $q \in \text{conv}(F \cap \bar{r})$ lies between boundary facet F and F_0 , let q_0 be the intersection of the line segment $q\bar{r}$ with F_0 ; then $\|q - q_0\| \leq d_0 = \delta$.
- (G2) If $q \in \text{conv}(F \cap \bar{r})$ lies between F_{i-1} and F_i and the segment $q\bar{r}$ intersects F_i at q_i , let q_F be the intersection of F with the ray $\bar{r}\bar{q}$; then $\|q_i - q\|/\|q_i - q_F\| \leq \rho = \delta/\varepsilon$.

We now describe how to handle certain simple types of regions: where f is a point or an edge. These will be handled the same regardless of the dimension of the problem, and they (the edge case in particular) will be used as important base cases for higher dimensional problems.

Point regions. Consider a point region $\sigma(p)$. For each $h \in \mathbb{H}$ create θ -net $U_{p,h}$ for $\Psi_{p,h}$, so $N_{p,h}$ are the corresponding points where each $p_i = p[h, u_i] \in N_{p,h}$ has $u_i \in U_{p,h}$. Put each $N_{p,h}$ in \mathcal{K}' .

For any point $q = \bar{q}[u', z] \in \sigma(p) \cap \text{conv}(P)$, let $q' = \bar{q}[u, h]$ where $h \in \mathbb{H}$ is the largest value such that $h \leq z$ and $u \in U_{p,h}$ is the closest direction to u' ; set $q_h = \bar{q}[u', h] = q'[u']$. First $\|q - q_h\| \leq \delta$ because $z - h \leq \delta$. Second $\|q_h - q'\| \leq \delta$ because the angle between u' and u is at most θ , and they are rotated about the point p . Thus $\|q - q'\| \leq \|q - q_h\| + \|q_h - q'\| \leq 2\delta \leq \varepsilon/2$.



Lemma 3.1. For a point region $\sigma(p)$, there exists a constant number of points $K_p \subset \sigma(p) \cap \text{conv}(P)$ such that all points $q \in \sigma(p) \cap \text{conv}(P)$ are within a distance $\varepsilon/2$ of $\text{conv}(K_p)$.

Edge regions. Consider an edge region $\sigma(e)$ for an edge e of \mathcal{P} . Orient e along the x -axis. For each $h \in \mathbb{H}$ and $u \in U_{e,h}$, let $\Psi_{e,h,u}$ be the set of points in $\Psi_{e,h}$ within an angle θ of u . For each $\Psi_{e,h,u}$, we add to K_e the (two) points of $\Psi_{e,h,u}$ with the largest and smallest x -coordinates, denoted by $p_{h,u}^+$ and $p_{h,u}^-$.





Figure 3. Illustration of 2 points in \mathcal{K}' for edge case with specific $h \in \mathbb{H}$ and $u \in U_{e,h,\theta}$.

For any point $q = \bar{q}[v, z] \in \sigma(e) \cap \text{conv}(P)$, there is a point $q'' = \bar{q}[u, h]$ such that $h \in \mathbb{H}$ is the largest value less than z and $u \in U_{e,h}$ is the closest direction to v . Furthermore, $\|q - q''\| \leq \|q - q_h\| + \|q_h - q''\| \leq (z - h) + 2\varepsilon \sin(\theta/2) = \delta + \delta = 2\delta$. We can also argue that there is a point $q' = \bar{q}[u', z'] \in p_{h,u}^- p_{h,u}^+$, because if \bar{q} has smaller x -coordinate than $p_{h,u}^-$ or larger x -coordinate than $p_{h,u}^+$, then q' cannot be in $\Psi_{e,h,u}$. Clearly the angle between u and u' is less than θ . This also implies that $h - z' < \delta$. Thus $\|q'' - q'\| \leq 2\delta$, implying $\|q - q'\| \leq 4\delta \leq \varepsilon/2$.

Lemma 3.2. For an edge region $\sigma(e)$, there exists $O(1)$ points $K_e \subset \sigma(e) \cap \text{conv}(P)$ such that for any point $q = \bar{q}[z, v] \in \sigma(e) \cap \text{conv}(P)$ there is a point $p = \bar{q}[h, u] \in \text{conv}(K_e)$ such that $z - h \leq 2\delta$, $\|v - u\| \leq 2\delta$, and, in particular, $\|q - p\| \leq 4\delta \leq \varepsilon/2$.

3.2.1 Approximation Stability in \mathbb{R}^2

For $\mathcal{K} \subset P \in \mathbb{R}^2$ there are $|\mathcal{K}|$ points and edges in \mathcal{P} . Thus combining Lemmas 3.1 and 3.2 $|\mathcal{K}'|/|\mathcal{K}| = O(1)$ and we have proven Theorem 1.2 for $d = 2$.

Theorem 3.1. For any point set $P \in \mathbb{R}^2$ and for any $\varepsilon > 0$ we have $\kappa(P, \varepsilon/2)/\kappa(P, \varepsilon) = O(1)$.

3.2.2 Approximation Stability in \mathbb{R}^3

Construction of \mathcal{K}' . Now consider $\mathcal{K} \subset P \in \mathbb{R}^3$ and the point regions, edge regions, and triangle regions in the decomposition of $\mathcal{P}_\varepsilon \setminus \text{intr } \mathcal{P}$ (see Figure 4). By Lemmas 3.1 and 3.2 we can add $O(|\mathcal{K}|)$ points to \mathcal{K}' to account for all point and edge regions. We can now focus on the $O(|\mathcal{K}|)$ triangle regions.




Figure 4. Illustration of regions in the partition of $\mathcal{P}_\varepsilon \setminus \text{intr } \mathcal{P}$ in three dimensions.

Consider a triangle region $\sigma(t)$ for a triangle t in \mathcal{P} (see Figure 5(a)), t^* consists of a single direction, the one normal to t . Let r be the highest point of $\sigma(t) \cap \text{conv}(P)$ in direction t^* . We add r to \mathcal{K}' and we create an exponential grid $G_{r,t}$ with r as the support point. For each edge $e \in G_{r,t}$ and $h \in \mathbb{H}$ we add the intersection of $e[t^*, h]$ with the boundary of $\sigma(t) \cap \text{conv}(P)$ to \mathcal{K}' , as shown in Figure 5(b). Thus, in total we add $O(|\mathcal{K}| \log(1/\varepsilon))$ points to \mathcal{K}' .




Figure 5. Illustration to aid correctness of approximation of triangle regions in \mathbb{R}^3 .

Proof of correctness. Consider any point $q = \bar{q}[t^*, z] \in \sigma(t) \cap \text{conv}(P)$ and associate it with a boundary edge e of t such that $\bar{q} \in \text{conv}(e \cup \bar{r})$. Let $q_h = \bar{q}[t^*, h]$ where $h \in \mathbb{H}$ is the largest height such that $h \leq z$. If segment $\bar{q}\bar{r}$ does not intersect any edge e_i parallel to e in $G_{r,t}$, let $\bar{p} = \bar{r}$. Otherwise, let e_i be the first segment parallel to e in $G_{r,t}$ intersected by the ray $\bar{q}\bar{r}$, and let \bar{p} be the intersection. Let $p = \bar{p}[t^*, h]$ which must be in $\text{conv}(\mathcal{K}')$ by construction. If $e_i = e_0$, then by (G1) we have $\|q_h - p\| = \|\bar{q} - \bar{p}\| \leq \delta$, thus $\|q - p\| \leq 2\delta \leq \varepsilon/2$ and we are done. Otherwise, let \bar{q}_e be the intersection of e with ray $\bar{r}\bar{q}$. By (G2) $\|\bar{p} - \bar{q}\|/\|\bar{p} - \bar{q}_e\| \leq \rho = \delta/\varepsilon$. Thus, $q'' = \bar{q}[t^*, h - \varepsilon\rho]$ is below the segment $\bar{q}_e p$ (see Figure 5(c)) and thus $q'' \in \text{conv}(\mathcal{K}')$ since triangle $p\bar{p}\bar{q}_e$ is in $\text{conv}(\mathcal{K}')$. Finally, $\|q - q''\| = \|q - q_h\| + \|q_h - q''\| \leq 2\delta \leq \varepsilon/2$. This proves Theorem 1.2 for $d = 3$.

Theorem 3.2. For any point set $P \in \mathbb{R}^3$ and for any $\varepsilon > 0$ we have $\kappa(P, \varepsilon/2)/\kappa(P, \varepsilon) = O(\log 1/\varepsilon)$.

3.2.3 Approximation Stability in \mathbb{R}^d

Construction of \mathcal{K}' . The number of regions in the decomposition of $\mathcal{P}_\varepsilon \setminus \text{intr } \mathcal{P}$ is $|\mathcal{P}| = O(|\mathcal{K}|^{\lfloor d/2 \rfloor})$ [14]. For each region $\sigma(f)$, we choose a set K_f of $O(\log^{d-2}(1/\varepsilon))$ points such that $\text{conv}(K_f) \subseteq \text{conv}(P) \cap \sigma(f) \subseteq \text{conv}(K_f) \oplus (\varepsilon/2)\mathbb{B}^d$. Then we set $\mathcal{K}' = \bigcup_f K_f$.

When f is a 0- or 1-simplex, we apply Lemmas 3.1 and 3.2. Otherwise, to construct K_f we create a recursive exponential grid on f . Specifically, for all $h \in \mathbb{H}$ and $u \in U_{f,h}$ we choose a point $r = \bar{r}[u, h] \in \sigma(f) \cap \text{conv}(P)$ and construct an exponential grid $G_{r,f}$ with r as the support point. Next, for each $\Delta \in G_{r,f}$ we recursively construct exponential grids. That is, for all $h' \in \mathbb{H}$ and $u' \in U_{\Delta,h'}$ we choose another support point $r' = \bar{r}'[u', h']$ and construct an exponential grid $G_{r',\Delta}$ on Δ . At each iteration the dimension of the simplex in the exponential grids drops by one. We continue the recursion until we get 1-simplices. Let \mathcal{G}_f be the union of all exponential grids.

Let e be a 1-simplex in \mathcal{G}_f . For each height $h \in \mathbb{H}$ and direction $u \in U_{e,h}$ we choose two points as described in the construction the edge region and add them to K_f . We also place the support point of each $\Delta \in \mathcal{G}_f$ into K_f . By construction, for a k -simplex f , \mathcal{G}_f contains $O(\log^{k-1}(1/\varepsilon))$ simplices and thus $|K_f| = O(\log^{k-1}(1/\varepsilon))$. Hence $|\mathcal{K}'| = O(|\mathcal{K}|^{\lfloor d/2 \rfloor} \log^{d-2}(1/\varepsilon))$.

Proof of correctness. Let f be a k -face of \mathcal{P} ($0 \leq k \leq d-1$). We need to show for any point $q = \bar{q}[u, z] \in \text{conv}(P) \cap \sigma(f)$, there is a point $p \in \text{conv}(K_f)$ (specifically p_k'') such that $\|q - p\| \leq \varepsilon/2$.

Before describing the technical details (mainly left to the appendix), we first provide some intuition regarding the proof. For any $q \in \text{conv}(P) \cap \sigma(f)$ we first consider $q_k = \bar{q}[u, h]$ where $h \in \mathbb{H}$ is the largest $h < z$. If $q_k \in \text{conv}(K_f)$ we are done. If $q_k \notin \text{conv}(K_f)$, we need to find a “helper point” q_{k-1} for q_k . If $q_{k-1} \notin \text{conv}(K_f)$ we need to recursively find a “helper point” q_{k-2} for q_{k-1} , and so on until $q_m \in \text{conv}(K_f)$. Tracing back along the recursion, we can then prove that q_k (and hence q) has a nearby point in $\text{conv}(K_f)$. Note that we do not prove q_j is near q_{j-1} . Formally:

Lemma 3.3. *We can construct a sequence of helper points $q_j = \bar{q}_j[u_j, h_j]$ and simplices $\Delta_j \in \mathcal{G}_f$ with the following invariants: We can construct a sequence of helper points $q_j = \bar{q}_j[u_j, h_j]$ and simplices $\Delta_j \in \mathcal{G}_f$ with the following invariants:*

- (I1) $q_j \in \text{conv}(P)$;
- (I2) $h_j \in \mathbb{H}$ and $h_j - h_{j+1} \leq \delta$ (for $j \neq k$);
- (I3) $\bar{q}_j \in \Delta_j \in \mathcal{G}_f$ and the dimension of Δ_j is j ; and
- (I4) $\|u_j - u_{j+1}\| \leq 2\theta$ (for $j \neq k$).

Proof. Set q_k as above and $\Delta_k = f$. Assume that for an index $j \leq k$, q_j has been defined. Since (I1) implies $q_j \in \Psi_{\Delta_j, h} \neq \emptyset$, there is a direction $u'_{j-1} \in U_{f, h_j}$ such that the angle between u_j and u'_{j-1} is at most θ . Let $r_j \in K_f$ be the support point of Δ_j , given h_j and u'_{j-1} . Let F be the facet of Δ_j such that $\bar{q}_j \in \text{conv}(F \cup \bar{r}_j)$. If the segment $\bar{q}_j \bar{r}_j$ does not intersect any simplex in the family of $(j-1)$ simplices induced by F and r_j , then let $q_{j-1} = r_j$ and terminate the recursion (see Figure 6(a)). Otherwise, let Δ_{j-1} be the first such simplex intersected by the ray $\overrightarrow{\bar{q}_j \bar{r}_j}$, and let \bar{q}_{j-1} be the intersection point (see Figure 6(b)).

To determine q_{j-1} we first find the direction u_{j-1} , such that $q'_{j-1} = \bar{q}_{j-1}[u_{j-1}, z']$ lies on the segment $q_j r_j$. Then h_{j-1} is determined as the maximum $h \in \mathbb{H}$ such that $h \leq z'$. We can show (I1) is satisfied because q'_{j-1} must be in $\text{conv}(P)$ because q_j and r_j are, and then $q_{j-1} = \bar{q}_{j-1}[u_{j-1}, h_{j-1}]$ lies on segment $q'_{j-1} \bar{q}'_{j-1}$. We show (I4) by $\|u_j - u_{j-1}\| \leq \|u_j - u'_{j-1}\| \leq \theta$. Invariant (I2) follows because $z' \geq h_j \cos \theta$ and thus $h_j - z' \leq h_j(1 - \cos \theta) \leq \varepsilon(\sin \theta) \leq \delta$. Invariant (I3) holds by construction.

Assume the recursion does not terminate early. At $j = 1$, since $q_1 \in \text{conv}(P)$ and Δ_1 is a line segment, we can apply Lemma 3.2 to $\sigma(\Delta_1)$ and find a point $q_0 = \bar{q}_1[u_0, h_1] \in \text{conv}(K_f)$ such that $\|q_1 - q_0\| \leq 2\delta$ and $\|u_1 - u_0\| \leq 2\theta$. This completes the description of the recursive definition of helper points. \square




Figure 6. (a) Terminate early: $q_{j-1} = r_j$. (b) Recursive case: q'_{j-1} lies on the segment $q_j r_j$.

Let q_m be the last point defined in the sequence. By construction, $q_m \in \text{conv}(K_f)$. For each $m \leq j \leq k$, let $p_j = q_j[u_m]$. We have the following key lemma, which shows that p_j is close to a point $p''_j \in \text{conv}(K_f)$.

Lemma 3.4. *For each $j \geq m$, there is a point $p'_j \in \Delta_j[u_m, h_j]$ such that*

- (1) $\|p'_j - p_j\| \leq j\delta$; and
- (2) $p''_j = p'_j - 2(j - m)\delta u_m \in \text{conv}(K_f)$.

Proof. We prove the lemma by induction on j . For $j = m$, since $p_m = q_m \in \text{conv}(K_f)$, the claim is trivially true by setting $p'_m = p_m$. Assume the claim is true for some $j \geq m$. Now consider the case $j + 1$. Let $\bar{p}_{F,j}$ be the intersection of the ray $\bar{q}_j \bar{q}_{j+1}$ with ∂f on facet F . Let x_{j+1} be the intersection of $\bar{p}_{F,j} \bar{p}'_j$ with the line passing through \bar{q}_{j+1} and parallel to $\bar{q}_j \bar{p}'_j$ (see Figure 7). There are two cases:




Figure 7. The inductive step for proving $\|p'_{j+1} - p_{j+1}\| \leq (j + 1)\delta$ and $p''_{j+1} \in \text{conv}(K_f)$.

Case 1: If Δ_j (such that $\bar{q}_j \in \Delta_j$ according to (I3)) is the closest facet to F , then q_{j+1} lies between F and Δ_j . Thus by (G1), we know that $\|\bar{q}_{j+1} - \bar{q}_j\| \leq \delta$. We set $p'_{j+1} = \bar{p}'_j[u_m, h_{j+1}]$. As such,

$$\|p'_{j+1} - p_{j+1}\| = \|\bar{p}'_j - \bar{q}_{j+1}\| \leq \|\bar{p}'_j - \bar{q}_j\| + \|\bar{q}_j - \bar{q}_{j+1}\| \leq j\delta + \delta = (j + 1)\delta.$$

Moreover, since $h_{j+1} - 2(j + 1 - m)\delta \leq h_j - 2(j - m)\delta$, p''_{j+1} lies on the segment $\bar{p}'_j \bar{p}'_j$ and therefore $p''_{j+1} \in \text{conv}(K_f)$.

Case 2: Otherwise. In this case, we have by (G2)

$$\frac{\|\bar{p}'_j - x_{j+1}\|}{\|\bar{p}'_j - \bar{p}_{F,j}\|} = \frac{\|\bar{q}_j - \bar{q}_{j+1}\|}{\|\bar{q}_j - \bar{p}_{F,j}\|} \leq \rho.$$

We set $p'_{j+1} = x_{j+1}[u_m, h_{j+1}]$. First observe that

$$\|p'_{j+1} - p_{j+1}\| = \|x_{j+1} - \bar{q}_{j+1}\| \leq \|\bar{p}'_j - \bar{q}_j\| \leq \delta.$$

Furthermore, let y be the intersection of $\bar{p}_{F,j}p''_j$ with $p'_{j+1}x_{j+1}$ (see Figure 7(right)). We then have

$$\|p'_{j+1} - y\| \leq (h_{j+1} - h_j) + \|p'_j - p''_j\| + \rho \cdot \|p''_j - \bar{p}'_j\| < \delta + 2(j - m)\delta + \delta = 2(j + 1 - m)\delta.$$

Therefore p''_{j+1} lies below y and as such $p''_{j+1} \in \text{conv}(K_f)$ since triangle $\bar{p}'_j \bar{p}_{F,j} p''_j \in \text{conv}(K_f)$. \square

We now complete the proof of Theorem 1.2 with the following lemma.

Lemma 3.5. $\|q - p''_k\| \leq \varepsilon/2$.

Proof. For $q = \bar{q}[u, z]$ and $q_k = \bar{q}[u, h]$, since $h \leq z \leq h + \delta$ then $\|q - q_k\| \leq \delta$.

Since $p_k = q_j[u_m]$, then invariant (I4) implies that $\|u_m - u_k\| \leq (k - m)2\theta \leq d2\theta$ and hence $\|p_k - q_k\| \leq 2\varepsilon \sin((1/2)d2\theta) \leq \varepsilon 2(d\delta/\varepsilon) \leq 2d\delta$.

Finally, for $j = k$, Lemma 3.4 implies that $\|p'_k - p_k\| \leq k\delta \leq d\delta$, that $\|p'_k - p''_k\| \leq 2(k - m)\delta \leq 2d\delta$, and $p''_k \in \text{conv}(K_f)$.

It follows that

$$\begin{aligned} \|q - p''_k\| &\leq \|q - q_k\| + \|q_k - p_k\| + \|p_k - p'_k\| + \|p'_k - p''_k\| \\ &\leq \delta + 2d\delta + d\delta + 2d\delta \\ &\leq 6d\delta = \varepsilon/2, \end{aligned}$$

as desired. \square

3.3 Remarks

- (1) For $d = 2, 3$, $\kappa(P, \varepsilon/2)$ is only a factor of $O(1)$ and $O(\log(1/\varepsilon))$, respectively, larger than $\kappa(P, \varepsilon)$; therefore, the sizes of optimal ε -kernels in these dimensions are relative stable. However, for $d \geq 4$, the stability drastically reduces in the worst case because of the superlinear dependency on $\kappa(P, \varepsilon)$.
- (2) Neither the upper nor the lower bound in the theorem is tight. For $d = 3$, we can prove a tighter lower bound of $\Omega(\kappa(P, \varepsilon) \log(1/(\varepsilon \cdot \kappa(P, \varepsilon))))$. We conjecture in \mathbb{R}^d that

$$\kappa(P, \varepsilon/2) = \Theta(\kappa(P, \varepsilon)^{\lfloor d/2 \rfloor} \log^{d-2}(1/(\varepsilon^{(d-1)/2} \cdot \kappa(P, \varepsilon)))).$$

References

- [1] P. K. Agarwal, S. Har-Peled, and K. Varadarajan, Approximating extent measure of points, *Journal of ACM*, 51 (2004), 606–635.
- [2] P. K. Agarwal, S. Har-Peled, and K. Varadarajan, Geometric approximations via coresets, in: *Combinatorial and Computational Geometry*, 2005, pp. 1–31.

- [3] P. K. Agarwal and H. Yu, A space-optimal data-stream algorithm for coresets in the plane, *SoCG*, 2007, pp. 1–10.
- [4] A. Bagchi, A. Chaudhary, D. Eppstein, and M. Goodrich, Deterministic sampling and range counting on geometric data streams, *SoCG*, 2004, pp. 144–151.
- [5] G. Barequet and S. Har-Peled, Efficiently approximating the minimum-volume bounding box of a point set in three dimensions, *Journal of Algs*, 38 (2001), 91–109.
- [6] T. Chan, Faster core-set constructions and data-stream algorithms in fixed dimensions, *Computational Geometry: Theory and Applications*, 35 (2006), 20–35.
- [7] T. Chan, Dynamic coresets, *SoCG*, 2008, pp. 1–9.
- [8] S. Har-Peled, *Approximation Algorithm in Geometry*, <http://valis.cs.uiuc.edu/~sariel/teach/notes/aprx/>, 2009.
- [9] J. Hershberger and S. Suri, Adaptive sampling for geometric problems over data streams, *Computational Geometry: Theory and Applications*, 39 (2008), 191–208.
- [10] M. H. Overmars, *The Design of Dynamic Data Structures*, Springer-Verlag, 1983.
- [11] S. Suri, C. D. Tóth, and Y. Zhou, Range counting over multidimensional data streams, *SoCG*, 2004, pp. 160–169.
- [12] H. Yu, P. K. Agarwal, R. Poredy, and K. Varadarajan, Practical methods for shape fitting and kinetic data structures using coresets, *Algorithmica*, 52 (2008).
- [13] H. Zarrabi-Zadeh, An almost space-optimal streaming algorithm for coresets in fixed dimensions, *ESA*, 2008, pp. 817–829.
- [14] G. M. Ziegler, *Lectures on Polytopes*, Springer-Verlag, 1995.